International Scientific Journal of Engineering and Management (ISJEM)
Volume: 04 Issue: 06 | June – 2025
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata
ISSN: 2583-6129
DOI: 10.55041/ISJEM03992

Reinforcement Learning in Business Decision Support Systems: Real-Time Optimization of Pricing and Inventory

Akshat Bajpai, Dr. Shilpa Pandey

[1] *Akshat Bajpai, Final Year Student, Amity Business School,*
Amity University Chhattisgarh, Raipur, India
[2] *Dr. Shilpa Pandey, , Assistant Professor, Amity Business School,*
Amity University Chhattisgarh, Raipur, India

---***---

Abstract - Amid a rapidly evolving market environment, enterprises must consistently recalibrate their pricing and inventory strategies to stay competitive and meet customer needs. Conventional Decision Support Systems (DSS) often depend on fixed models and rules, rendering them ineffective in dynamic scenarios. This paper introduces a smart DSS architecture that leverages Reinforcement Learning (RL), focusing on Q-Learning and Deep Q-Networks (DQN), to enable real-time optimization of pricing and inventory decisions. Through simulations in a retail setting characterized by fluctuating demand and uncertain supply chains, the RL agent is trained to boost overall profit while curbing stockouts and inventory costs. Our findings show substantial performance gains over traditional rule-based models, with adaptive strategies developing through learned behavior. The study underscores RL's capacity to drive autonomous and data-centric improvements in business decision-making.

***Key Words*:** Reinforcement Learning, DSS, Pricing Strategy, Inventory Optimization, Q-Learning, DQN, Real-Time Decisions

1. INTRODUCTION

1.1 Background

Modern supply chains and customer behavior evolve rapidly, creating the need for intelligent and responsive decision-making in business operations. Key areas such as pricing and inventory are still often managed with outdated, static techniques.

1.2 Research Problem

Traditional DSS approaches often rely on predictive models and static rules, making them less effective when confronted with real-time changes in consumer demand or supply uncertainties. Reinforcement Learning presents a viable solution, enabling systems to learn optimal responses through interaction with the environment.

1.3 Research Goals

- Develop a simulated environment to reflect dynamic business operations.
- Apply and test RL algorithms like Q-Learning and DQN for decision support.
- Measure outcomes using metrics like profitability, stock turnover, and service level.
- Compare the RL-enhanced DSS against traditional decision-making frameworks.

2. Literature Review

2.1 Decision Support Systems

DSS are computer-based systems intended to assist in business decision-making. Although effective in structured settings, traditional systems lack adaptive capabilities unless manually updated.

2.2 Pricing and Inventory Control

Pricing: Conventional pricing strategies include cost-plus and elasticity-based dynamic pricing.

Inventory: Techniques such as EOQ, ABC, and Just-in-Time dominate but assume static market conditions.

2.3 Role of RL in Business

Although widely applied in areas like robotics and automated control systems, RL methods (especially model-free ones like Q-Learning) are still underutilized in business contexts. However, current research indicates potential for enhancing pricing and supply chain resilience.

ISSN: 2583-6129
DOI: 10.55041/ISJEM03992

3. Methodology

3.1 System Overview

Key Components:

- Agent: The RL-powered decision support mechanism.
- Environment: Simulated retail setup with varying demand.
- States: Includes factors like inventory status, time, price tier, and past demand.
- Actions: Price adjustments and inventory replenishments.
- Rewards: Calculated based on profit, inventory holding costs, and stockout penalties.

3.2 Algorithms

- Q-Learning: A table-based reinforcement learning technique that updates an action-value function using the Bellman equation.
- DQN: An advanced RL method using neural networks to estimate Q-values, supported by experience replay and target networks.

4. Experimental Design

4.1 Simulation Parameters

Parameter	Value
Time Frame	10,000 episodes
Demand Distribution	Poisson (λ varies)
Max Inventory	100 units
Holding Cost	$1/unit/step
Stockout Penalty	$5/unit
Price Range	$10 - $50

4.2 Tools and Technologies

- Python 3.9
- Custom OpenAI Gym environment
- TensorFlow / Keras for deep learning
- Matplotlib and Seaborn for data visualization

5. Results

5.1 Agent Training

Over 6,000 episodes, cumulative rewards steadily increased, indicating learning and performance improvement.

5.2 Inventory Efficiency

Compared to rule-based methods, the RL agent maintained more stable inventory levels and minimized stockouts.

5.3 Pricing Trends

RL-driven pricing adjusted dynamically based on demand patterns and inventory status.

5.4 Performance Metrics

Metric	Rule-Based	RL (DQN)
Profit ($)	18,320	22,430
Stockouts (#)	140	90
Avg Inventory	65	50
Service Level (%)	86	95

6. Discussion

6.1 Findings

- RL-based systems demonstrated superior adaptability to fluctuating demand.
- Better inventory control resulted in cost savings and improved customer service.
- Pricing policies evolved to resemble real-world dynamic pricing models.

6.2 Constraints

- The study is simulation-based and not yet validated in live business scenarios.
- RL requires an initial exploration phase, which can be inefficient.
- The decision logic of deep RL models can be difficult to interpret.

ISSN: 2583-6129
DOI: 10.55041/ISJEM03992

7. Conclusion

Integrating RL into business decision systems shows notable potential for enhancing real-time decision-making in pricing and inventory management. The results indicate a clear advantage over conventional models. Future work may explore multi-agent interactions, deployment in real-world businesses, and combining RL with expert-based rules.

References

1. Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction.

2. Van Roy, B., & Tsitsiklis, J. N. (1999). Q-learning with Function Approximation.

3. Silver, D., et al. (2016). Mastering the Game of Go with Deep Neural Networks. Nature.

4. OpenAI Gym Documentation.

5. Bertsekas, D. P. (2007). Dynamic Programming and Optimal Control.